

CLAIMS

What is claimed is:

1. A method for per-session network address translation (NAT) learning in a media gateway, the method comprising:
 - 5 in a media gateway:
 - (a) receiving a media session setup request for establishing a media session;
 - (b) in response to the media session setup request, assigning a local network and transport address combination identifying a media processing resource within the media gateway for processing a media stream associated with the media session;
 - 10 (c) receiving at least one initial media packet in the media stream, the initial media packet being addressed to the local network and transport address combination and having a source network address and a source transport address, at least one of the source network address and the source transport address being assigned by a network address translator;
 - 15 (d) learning the source network address from the initial media packet;
 - 20 (e) processing the initial media packet using the media processing resource assigned to the session;
 - (f) accepting and processing subsequent media packets having the assigned local network address and local transport address in

their destination address fields and the learned source network address in their source address fields; and

(g) repeating steps (a)-(f) for each new incoming session to the media gateway and thereby performing NAT learning on a per-session basis.

5

2. The method of claim 1 wherein receiving a media session setup request includes receiving a request from a soft switch to allocate resources for a new media session.

10

3. The method of claim 1 wherein the media session comprises at least one voice call.

4. The method of claim 1 wherein the media stream comprises a Real-time Transmission Protocol (RTP) media stream.

15

5. The method of claim 1 wherein assigning a local network and transport address combination includes assigning the local network and transport address combination to a media processing chip for processing the media stream.

6. The method of claim 1 wherein learning the source network address includes:

20

(a) receiving the initial media packet at the media processing resource;

(b) routing the initial media packet from the media processing resource to a central processing unit (CPU) operatively associated with the media processing resource; and

(c) at the CPU, extracting the source network address and from the initial media packet and broadcasting the learned source network address to a plurality of network interface cards in the media gateway.

- 5 7. The method of claim 6 comprising learning the source transport address from the initial media packet and broadcasting the source transport address to the plurality of network interface cards in the media gateway.
8. The method of claim 7 comprising, at the network interface cards, using the learned source network address, the learned source transport address, the local network address, and the local transport address to
10 create a per-session pin-hole for firewall filtering.
9. The method of claim 1 wherein learning the source network address includes dynamically assigning one of a plurality of distributed media processing elements in the media gateway to learn the source network
15 address.
10. The method of claim 1 wherein learning the source network address includes:
- (a) receiving the initial media packet at the media processing resource; and
- 20 (b) at the media processing resource, extracting the source network address and broadcasting the learned source network address to a plurality of network interface cards in the media gateway.
11. The method of claim 10 comprising learning the source transport address from the initial media packet at the media processing resource

and broadcasting the learned source transport address to the plurality of network interface cards in the media gateway.

12. The method of claim 11 comprising, at the network interface cards, using the learned source network address, the learned source transport address, the local network address, and the local transport address to create a per-session pin-hole for firewall filtering.
13. The method of claim 10 wherein the media stream comprises a voice-over-IP-to-voice-over-IP media stream and wherein accepting and processing subsequent media packets for the session includes receiving subsequent media packet associated with the session at a first network interface card, determining a destination network interface card based on the destination address, and forwarding all the subsequent media packets to the selected destination network interface card.
14. The method of claim 1 comprising, after step (d), performing firewall filtering for the subsequent media packets using the local network address, the local transport address, the source network address, and the source transport address.
15. The method of claim 14 wherein performing firewall filtering includes rejecting media packets that have the local network address and the local transport address in their destination address fields but do not have the source network address and the source transport address in their source address fields.
16. The method of claim 1 wherein the media session comprises a voice call and wherein the method further comprises seamlessly inserting an

internal media processor into the call without changing topology of the call during any time of the call, including call initialization time, call active state, and call release time.

17. The method of claim 16 wherein inserting an internal media processor
5 into the call includes inserting at least one of: an announcement server, a conference bridge, a DTMF generator, a DTMF collector, a voice mail server, and a law enforcement circuit into the call.
18. The method of claim 1 wherein the media session comprises a voice call and wherein the method further comprises comprising seamlessly
10 inserting an external media processor into the call without changing topology of the call for the duration of the call, including call initialization time, call active state, and call release time.
19. The method of claim 18 wherein inserting an internal media processor into the call includes inserting at least one of: an announcement server,
15 a conference bridge, a DTMF generator, a DTMF collector, a voice mail server, and a law enforcement circuit into the call.
20. A media gateway with internal network address translation (NAT) learning capabilities, the media gateway comprising:
- 20 (a) a plurality of network interface cards for receiving media packets, for determining whether the media packets have been assigned to a session, and for forwarding the media packets that have been assigned to a session to a media processing resource;
- (b) a plurality of media processing resources for processing the media packets that have been assigned to a session; and

- (c) a NAT learning function located within the media gateway and operatively associated with the media processing resources and the network interface cards for learning dynamically assigned source addresses assigned to media packets and for communicating the learned source addresses to the network interface cards.
- 5
21. The media gateway of claim 20 wherein the network interface cards comprise packet network interface cards.
22. The method of claim 20 wherein the network interface cards comprise
- 10 ATM network interface cards.
23. The media gateway of claim 20 wherein the media processing resources include voice-over-IP SAR chips for processing voice-over-IP calls.
24. The media gateway of claim 23 wherein the NAT-learning function is performed by the voice-over-IP SAR chips.
- 15 25. The media gateway of claim 24 comprising a plurality of voice server modules associated with the voice-over-IP SAR chips, a central processing unit located on each voice server module for controlling the voice-over-IP SAR chips, wherein the NAT learning function is performed by one of the central processing units that is dynamically
- 20 assigned to the session.
26. The media gateway of claim 20 wherein the media processing resources include a first codec and a second codec and wherein the first codec and the second codec are used to perform transcoding for at least one

of voice-over-IP to voice-over-IP calls, voice-over-IP to voice-over-AAL1 calls and voice-over-IP to voice-over-AAL2 calls.

The media gateway of claim 20 wherein the NAT learning function is adapted to learn the source network address and the source transport address and to distribute the learned source network address and the learned source transport address to at least one of the network interface cards and wherein the network interface cards are adapted to accept media packets addressed to a local network address and local transport address assigned to the session and from the learned source network address and the learned source transport address.

The media gateway of claim 27 wherein the network interface cards are adapted to reject media packets addressed to the local source network address and local source transport address assigned to the session but that do not have the dynamically learned source network address and dynamically learned source transport address assigned to the session.

The media gateway of claim 20 wherein the NAT learning function is adapted to selectively filter media packets for each session based on a local network address, a local transport address, a dynamically learned source address, and a dynamically learned transport assigned to each session, thereby performing firewall filtering on a per-session basis.

A computer program product for per-session network address translation (NAT) learning in a media gateway, the computer program product comprising computer executable instructions embodied in a

computer readable medium for performing steps comprising:

in a media gateway:

- (a) receiving a media session setup request for establishing a media session;
- 5 (b) in response to the media session setup request, assigning a local network and transport address combination identifying a media processing resource within the media gateway for processing a media stream associated with the media session;
- 10 (c) receiving at least one initial media packet in the media stream, the initial media packet being addressed to the local network and transport address combination, the initial media packet having a source network address and a source transport address, at least one of the source network address and the source transport address being assigned by a network address translator;
- 15 (d) learning the source network address;
- (e) processing the initial media packet using the media processing resource assigned to the session;
- (f) accepting and processing subsequent media packets having the local network address and local transport address in their destination address fields and the learned source network address in their source address fields; and
- 20 (g) repeating steps (a)-(f) for each new session to the media gateway and thereby performing NAT learning on a per-session basis.

31. The computer program product of claim 30 wherein receiving a media session setup request includes receiving a request for allocating resources for a new media session from a soft switch.
32. The computer program product of claim 30 wherein the media session
5 comprises a voice call.
33. The computer program product of claim 30 wherein the media stream comprises a Real-time Transmission Protocol (RTP) media stream.
34. The computer program product of claim 30 wherein the media stream
10 comprises a Real-time Transmission Control Protocol (RTCP) media stream.
35. The computer program product of claim 30 wherein assigning a local network and transport address combination includes assigning the local network and transport address combination to a media processing chip for processing the media stream.
- 15 36. The computer program product of claim 30 wherein learning the source network address includes:
- (a) receiving the initial media packet at the media processing resource;
 - (b) routing the initial media packet from the media processing
20 resource to a central processing unit (CPU) operatively associated with the media processing resource; and
 - (c) at the CPU, extracting the source network address from the initial media packet and broadcasting the learned source network

address to a plurality of network interface cards in the media gateway.

37. The computer program product of claim 36 comprising learning the source transport address by extracting the source transport address from the initial media packet and broadcasting the source transport address to the plurality of network interface cards in the media gateway.
38. The computer program product of claim 37 comprising, at the network interface cards, using the learned source network address, the learned source transport address, the local network address, and the local transport address to create a per-session pin-hole for firewall filtering.
39. The computer program product of claim 30 wherein learning the source network address includes:
- (a) receiving the initial media packet at the media processing resource; and
 - (b) at the media processing resource, extracting the source network address from the initial media packet and broadcasting the learned source network address to a plurality of network interface cards in the media gateway.
40. The computer program product of claim 39 comprising learning the source transport address by extracting the source transport address from the media packet and broadcasting the learned source transport address to the plurality of network interface cards in the media gateway.
41. The computer program product of claim 40 comprising, at the network interface cards, using the learned source network address, the learned

source transport address, the local network address, and the local transport address to create a per-session pin-hole for firewall filtering.

42. The computer program product of claim 39 wherein the media stream comprises a voice-over-IP-to-voice-over-IP media stream and wherein
- 5 accepting and processing subsequent media packets includes receiving subsequent media packets associated with the session at the first network interface card, determining a destination network interface card based on a destination address in the subsequent media packets, and forwarding the subsequent media packets to the selected network
- 10 interface card.
43. The computer program product of claim 42 wherein accepting and processing subsequent media packets include performing transcoding for the media packets.
44. The computer program product of claim 42 wherein accepting and
- 15 processing subsequent media packets includes forwarding the subsequent media packets to the selected network interface card without performing transcoding.
45. The computer program product of claim 36 comprising, after step (c), performing firewall filtering for the subsequent media packets associated
- 20 with each session using the local network address, the local transport address, the learned source network address, and the learned source transport address.
46. The method of claim 45 wherein performing firewall filtering includes rejecting media packets that have the local network address and the

local transport address in their destination address fields but do not have the learned source network address and the learned source transport address in their source address fields.

47. The computer program product of claim 30 wherein the session
5 comprises a voice call and wherein the computer program product further performs the step of seamlessly inserting an internal media server into the call without changing topology of the call during any time of the call, including call initialization time, active state, and call release time.
- 10 48. The system of claim 47 wherein seamlessly inserting an internal media processor into the call includes seamlessly inserting at least one of: an announcement player, a conference bridge, a DTMF generator, a DTMF collector, a voice mail server, and a law enforcement circuit into the call.
49. The computer program product of claim 30 wherein the session
15 comprises a voice call and wherein the computer program product further performs the step of seamlessly inserting an external media processor into the call without changing topology of the call during any time of the call, including call initialization time, active state, and call release time.
- 20 50. The system of claim 49 wherein seamlessly inserting an external media processor into the call includes seamlessly at least one of: an announcement player, a conference bridge, a DTMF generator, a DTMF collector, a voice mail server, and a law enforcement circuit into the call.